Accurate quantum state estimation via "Keeping the experimentalist honest"

Robin Blume-Kohout
*Institute for Quantum Information, Caltech 107-81, Pasadena, CA 91125 USA**

Patrick Hayden
School of Computer Science, McGill University, Montreal, QC H3A 2A7 Canada†

In this article, we derive a unique procedure for quantum state estimation from a simple, self-evident principle: an experimentalist's estimate of the quantum state generated by an apparatus should be constrained by honesty. A skeptical observer should subject the estimate to a test that guarantees that a self-interested experimentalist will report the true state as accurately as possible. We also find a non-asymptotic, operational interpretation of the quantum relative entropy function.

arXiv:quant-ph/0603116v1 14 Mar 2006

Consider a source of quantum states such as a laser, or an ion trap with a preparation procedure. Quantum state estimation is the problem of deducing *what* state it emits by analyzing the outcomes of measurements on many instances. The usual procedure for state estimation is *quantum state tomography* [1, 2], together with some variant of *maximum-likelihood estimation* [3, 4] to ensure positivity. The obvious goal is an estimate "close" to the true state. Different metrics, such as fidelity [5], relative entropy [6], trace norm, or Hilbert-Schmidt norm [7], will favor different estimation procedures.

Here, we derive an optimal state-estimation procedure by first identifying quantum relative entropy as a unique metric for characterizing an estimate's "goodness". Our procedure is broadly adaptable to (1) arbitrary prior knowledge (or ignorance) and (2) arbitrary measurement procedures.

Keeping the Experimentalist Honest: Implicit in the idea of state estimation is the assumption that some estimates are better than others. Suppose that σ is an estimate of the true state ρ, and that $f(\rho:\sigma)$ is a measure of how "good" an estimate σ is. We propose that this measure should obey three principles:

1. The best estimate of ρ is ρ itself. If $f(\rho:\sigma)$ measures how well σ estimates ρ, then $f(\rho:\rho) > f(\rho:\sigma)$ for all $\sigma \neq \rho$.

2. $f(\rho:\sigma)$ should correspond to some *operational* test, as the payoff or cost of some experimental procedure.

3. The "reward" for correctly predicting an event should depend only on the predicted probability for *that* event. This is a version of the likelihood principle (see [8]).

Remarkably, these simple assumptions single out one measure: the *relative entropy* between ρ and σ, or $S(\rho\|\sigma) = \mathrm{Tr}(\rho\ln\rho - \rho\ln\sigma)$. It arises as the expected payoff in a type of game between a cash-strapped experimentalist and her employer.

Alice, an ambitious scientist attempting to build a quantum computer, produces states that she believes are described by the density operator ρ. She informs her employer, Bob, that she has produced the state σ. Bob, a conscientious scientific administrator, would like to ensure that Alice does not lie – that $\sigma = \rho$. He will periodically visit Alice's lab and measure one of her states, in a way that may depend on her estimate σ. Her future funding will depend on the outcomes of these measurements. What measurement should Bob perform, and how should he pay Alice, so that she has no incentive to deceive him?

We propose that Bob should measure in a basis $\{|f_i\rangle\}_{i=1}^n$ that diagonalizes $\sigma = \sum_i s_i|f_i\rangle\langle f_i|$. Upon getting outcome i, he should pay Alice $R_i = C + D\ln s_i$ dollars, where C and D are non-negative constants. We denote this as the "honest experimentalist reward scheme," or HERS.

HERS motivates honesty: Bob's measurement yields outcome i with probability $p_i = \mathrm{Tr}(\rho|f_i\rangle\langle f_i|)$. Alice's expected reward is

$$\overline{R} = \sum_{i=1}^n p_i R_i = C + D\sum_{i=1}^n p_i \ln s_i. \tag{1}$$

Rewriting the last term as $\sum_i p_i \ln s_i = \mathrm{Tr}\rho\ln\sigma$ yields

$$\begin{aligned}\overline{R}(\rho:\sigma) &= C + D\,\mathrm{Tr}\rho\ln\sigma \\ &= C + D\left[\mathrm{Tr}\rho\ln\rho - (\mathrm{Tr}\rho\ln\rho - \mathrm{Tr}\rho\ln\sigma)\right] \\ &= C - D\left[H(\rho) + S(\rho\|\sigma)\right] \end{aligned} \tag{2}$$

Since C, D, and ρ are fixed, Alice maximizes her expected reward by reporting a σ that minimizes the relative entropy $S(\rho\|\sigma)$. This constrains σ to be ρ itself [9]. Alice is thereby motivated to be honest. She is also motivated to produce pure states – but *not* to lie about how pure the actual state is.

HERS is unique: Unless Bob can do non-projective POVMs [31], this turns out to be the *only* verification procedure that satisfies our three criteria.

In classical statistics, a reward scheme for a probabilistic forecast is a *scoring rule*. It assigns an average reward $\overline{R}(P:Q)$ to a forecast Q when events are distributed according to P. A reward that is uniquely maximized by an

*Electronic address: robin@blumekohout.com
†Electronic address: patrick@cs.mcgill.ca

honest forecast is a *strictly proper scoring rule* or SPSR (see review [10]). Given some P, the maximum reward under such a rule is P's *value*, $G(P) \equiv \overline{R}(P : P)$. Savage showed that for every SPSR, $G(P)$ is strictly convex [11].

To consider the quantum case, we observe that a measurement transforms a state ρ into a probability distribution $\{P_i\}$ over outcomes, to which we can apply a scoring rule. We represent a projective measurement of basis B as a quantum channel β. For any state ρ, let $\beta[\rho] = P$. P is a diagonal matrix of probabilities, and β simply annihilates off-diagonal elements. Let $G(\rho)$ be the maximum of $G(\beta[\rho])$ over all β. Since the eigenvalues of ρ majorize those of $\beta[\rho]$ [12], and G is convex, this maximum is achieved when $\beta[\rho] = \rho$. Thus $G(\rho) = G(\{\lambda_i\})$, where $\{\lambda_i\}$ are the eigenvalues of ρ.

Lemma 1. *Given a physical state ρ and an estimate σ, Bob can ensure Alice's honesty by applying a SPSR to the probabilities for a measurement of basis B if and **only** if σ is diagonal in B.*

Proof: Represent Bob's σ-dependent measurement as a CP-map β_σ that annihilates off-diagonal elements in basis B. Let the SPSR yield a value G.

1. Suppose that B diagonalizes σ, so $\beta_\sigma[\sigma] = \sigma$. Then Alice's expected reward is

$$\overline{R}(\beta_\sigma[\rho] : \beta_\sigma[\sigma]) \leq G(\beta_\sigma[\rho]) \leq G(\rho). \tag{3}$$

The inequalities are simultaneously saturated if and only if $\sigma = \rho$, in which case Alice earns the full value $G(\rho)$ of her state. When $\sigma \neq \rho$, one or both of the inequalities is strict, so Alice earns strictly less than $G(\rho)$. Thus, Alice maximizes her reward uniquely by reporting $\sigma = \rho$.

2. Suppose that there exists a σ so that $\beta_\sigma[\sigma] \neq \sigma$. Then let $\rho = \beta_\sigma[\sigma]$. Since $\beta_\sigma^2 = \beta_\sigma$, $\beta_\sigma[\rho] = \rho = \beta_\sigma[\sigma]$. Alice's expected reward is

$$\overline{R}(\beta_\sigma[\rho] : \beta_\sigma[\sigma]) = \overline{R}(\rho, \rho) = G(\rho). \tag{4}$$

whereas if she (truthfully) reports ρ, she can expect

$$\overline{R}(\beta_\rho[\rho] : \beta_\rho[\rho]) = G(\beta_\rho[\rho]) \leq G(\rho), \tag{5}$$

where the inequality holds because ρ's eigenvalues majorize those of $\beta_\rho[\rho]$ (by Schur's theorem [12]), and because G is convex. Alice expects the same reward for reporting ρ or $\sigma \neq \rho$, so her honesty is not ensured. □

So far we have demanded only that our scoring rule be strictly proper. We now demand that Alice's reward depend only on her predicted probability for the *observed* event. In other words, $R_i(\{s_j\}) = R_i(s_i)$. This reflects the Likelihood Principle [13]: all the relevant information in an event is contained in the likelihood of the hypothesis (here, $p(i|\sigma)$). *How* the experiment was performed is irrelevant. In particular, this avoids any argument between Alice and Bob about how to describe the outcome[s] that did not occur.

A remarkable theorem by Aczel (see also [8]) then restricts the form of the reward function $R_i(s_i)$.

Theorem 1 (Aczel [14]). *Let $n \geq 3$. The inequality*

$$\sum_{i=1}^{n} p_i R_i(q_i) \leq \sum_{i=1}^{n} p_i R_i(p_i) \tag{6}$$

is satisfied for all n-point probability distributions $(p_1 \ldots p_n)$ and $(q_1 \ldots q_n)$ if and only if there exist constants $C_1 \ldots C_n$ and D such that for all $i \in [1 \ldots n]$,

$$R_i(p) = D \ln p + C_i. \tag{7}$$

In the scenario we consider, there is even less freedom. Aczel's theorem allows the constants C_i to depend on i. In the quantum setting, all the C_i must be equal to a fixed C *independent* of i (see proof in Appendix A). We have therefore proved the following:

Theorem 2 (The honest experimentalist). *Let A be a quantum system with dimension $n \geq 3$. Let ρ and σ be density operators for A, and let $\{|g_i\rangle\}$ be an orthonormal basis for A that depends only on σ. Defining $p_i = \langle g_i | \rho | g_i \rangle$ and $s_i = \langle g_i | \sigma | g_i \rangle$, suppose that*

$$\sum_i p_i R_i(s_i) \leq \sum_i p_i R_i(p_i) \tag{8}$$

is satisfied for all ρ and σ, with equality if and only if $\rho = \sigma$. Then $\{|g_i\rangle\}$ diagonalizes σ, and there exist constants C and D such that $R_i(s_i) = C + D \ln s_i$ for all i.

The scheme outlined (HERS) is *uniquely* specified by Bob's need to guarantee Alice's honesty through self-interest. Relative entropy, $S(\rho\|\sigma)$, appears naturally as the amount of money that Alice can expect to lose by lying. Any "boss" who wishes never to be lied to must use the HERS payment scheme.

Of course, rewards in the real world are often structured less wisely. We propose, however, that a maximally ethical scientist should *act* as if she were being motivated by HERS, and use $S(\rho\|\sigma)$ as the universal measure of *honesty*. Hereafter, we will assume that the experimentalist is, in fact, honest.

The Uncertain Experimentalist: What does "honesty" mean for an experimentalist who is not certain of ρ? We assert that she should behave *as if* she were guided by a tangible, strictly proper, reward scheme. HERS is an excellent candidate, but our proofs hold for any strictly proper scheme.

Suppose that Alice does not know ρ, but knows that it will be selected from an ensemble $\pi(\rho)\mathrm{d}\rho$ (or simply $\pi(\rho)$ hereafter, for clarity). Equivalently, she thinks the true state is ρ with probability $\pi(\rho)$. Her expected reward (from HERS) for reporting σ is:

$$\begin{aligned}
\overline{R} &= \int R(\rho : \sigma)\pi(\rho)\mathrm{d}\rho \\
&= C - D\left(\int H(\rho)\pi(\rho)\mathrm{d}\rho + \int S(\rho\|\sigma)\pi(\rho)\mathrm{d}\rho\right) \\
&= \mathrm{const}_\sigma + D\int \mathrm{Tr}(\rho\ln\sigma)\pi(\rho)\mathrm{d}\rho \\
&= \mathrm{const}_\sigma + D\left[\mathrm{Tr}(\overline{\rho}\ln\sigma)\right] = \mathrm{const}'_\sigma - D\left[S(\overline{\rho}\|\sigma)\right],
\end{aligned}$$

where $\overline{\rho} = \int \rho\pi(\rho)\mathrm{d}\rho$. The "const$_\sigma$" terms are independent of σ and therefore out of Alice's control. Therefore, Alice maximizes her honesty by reporting the mean of her probability distribution.

The uniqueness of HERS depends on the Likelihood Principle. However, the mean of the probability distribution is maximally honest for *any* strictly proper scoring rule:

Theorem 3. *Let Alice believe that ρ is selected from a distribution $\pi(\rho)$. Let her expected reward for reporting σ be $R(\rho : \sigma) = \sum_i p_i R_i(\sigma)$, where $p_i = \mathrm{Tr}E_i\rho$, and $R(\rho : \rho) > R(\rho : \sigma)$ for all $\sigma \neq \rho$. Alice maximizes her expected reward by reporting $\sigma = \overline{\rho} \equiv \int \rho\pi(\rho)\mathrm{d}\rho$.*

Proof: Since Alice expects ρ to appear with probability $\pi(\rho)$, her expected reward is:

$$\begin{aligned} \overline{R} &= \int R(\rho : \sigma)\pi(\rho)\mathrm{d}\rho && (9)\\ &= \int \sum_i \mathrm{Tr}(E_i\rho)R_i(\sigma)\pi(\rho)\mathrm{d}\rho && (10)\\ &= \sum_i \mathrm{Tr}E_i \left(\int \rho\pi(\rho)\mathrm{d}\rho\right) R_i(\sigma) && (11)\\ &= R(\overline{\rho} : \sigma). && (12)\end{aligned}$$

R is strictly proper, so the unique maximum of $R(\overline{\rho} : \sigma)$ is at $\sigma = \overline{\rho}$. □

Consider, instead if Alice had tried to maximize fidelity [15], which is not derived from an operational procedure, but *would* guarantee Alice's honesty when she knows ρ exactly [32]. Suppose that Alice knows that ρ is either $|0\rangle\langle 0|$ or $|+\rangle\langle +|$, with equal probability. The fidelity between any σ and a pure state is $F(\sigma, |\psi\rangle\langle\psi|) = \langle\psi|\,\sigma\,|\psi\rangle$, so the *average* fidelity is just:

$$\overline{F} = \mathrm{Tr}\left(\sigma\overline{\rho}\right), \qquad (13)$$

where $\overline{\rho} = \frac{1}{2}\left(|0\rangle\langle 0| + |+\rangle\langle +|\right)$. To maximize $\overline{F}$, Alice would choose the largest eigenstate of $\overline{\rho}$ – *not* $\overline{\rho}$ itself. Thus, while fidelity appears at first like a good measure of honesty, it does *not* generally motivate Alice to report the mean of her distribution. Moreoever, it can motivate her to report a pure state that she knows is not the true state.

This is not simply a different definition of honesty. An experimentalist who reports a pure state $|\psi\rangle$ is predicting that some event will *never* occur. If Alice reports $|0\rangle$, she is asserting that no measurement will ever yield $|1\rangle$. In the presence of any uncertainty whatsoever, this is at best misleading, and at worst an outright lie.

This illustrates that HERS strongly penalizes over-optimism. If Bob obtains an outcome $|f_i\rangle\langle f_i|$ for which $\langle f_i|\,\sigma\,|f_i\rangle = 0$, Alice will lose infinitely much money! A truly zero-probability event is one against which a gambler would bet infinite money, at arbitrarily bad odds. Reporting $p = 0$ for an event that could conceivably happen is infinitely misleading, and should be discouraged.

The Informed Experimentalist: How should Alice use the results of measurements (that she has performed) to reduce her uncertainty? Suppose that she has performed POVM measurements on N copies of ρ, where the ith result corresponds to a positive operator E_i. She knows two things:

1. ρ is selected at random from an ensemble described by $\pi_0(\rho)$.
2. Through experiments on copies of ρ, she has obtained a measurement record $\mathcal{M} = \{E_1, E_2 \ldots E_N\}$.

Suppose that she reports σ_j when $\mathcal{M}_j$ occurs, and is paid according to a SPSR where $\overline{R}(\rho : \sigma) = \sum_i \mathrm{Tr}(|f_i\rangle\langle f_i|\rho)R_i(\sigma)$.

Since ρ appears with probability $\pi_0(\rho)$, the event "ρ appeared, $\mathcal{M}_j$ was measured, and σ_j was reported" occurs with probability $\pi_0(\rho)p\left(\mathcal{M}_j|\rho\right)$. Alice's expected reward over *all* possible events is:

$$\begin{aligned} \overline{R} &= \int \sum_j R(\rho : \sigma_j)p\left(\mathcal{M}_j|\rho\right)\pi_0(\rho)\mathrm{d}\rho\\ &= \sum_j \int \left(\sum_i \mathrm{Tr}(|f_i\rangle\langle f_i|\rho)R_i(\sigma_j)\right) p\left(\mathcal{M}_j|\rho\right)\pi_0(\rho)\mathrm{d}\rho\\ &= \sum_{i,j} \mathrm{Tr}\left[|f_i\rangle\langle f_i| \left(\int \rho p\left(\mathcal{M}_j|\rho\right)\pi_0(\rho)\mathrm{d}\rho\right)\right] R_i(\sigma_j).\end{aligned}$$

We rewrite this using

$$\begin{aligned} p_j &\equiv \int p\left(\mathcal{M}_j|\rho\right)\pi_0(\rho)\mathrm{d}\rho, \text{ and} && (14)\\ \overline{\rho}_j &\equiv \frac{1}{p_j}\int \rho p\left(\mathcal{M}_j|\rho\right)\pi_0(\rho)\mathrm{d}\rho, && (15)\end{aligned}$$

to get

$$\overline{R} = \sum_j p_j \mathrm{Tr}\left(|f_i\rangle\langle f_i|\overline{\rho}_j\right) R_i(\sigma_j) = \sum_j p_j R(\overline{\rho}_j : \sigma_j).$$

$\overline{R}$ is strictly proper, so by setting $\sigma_j = \overline{\rho}_j$ we uniquely maximize each term in the sum, and Equation 15 defines the optimal estimate of ρ, given $\mathcal{M}_j$.

Equation 15 is nothing other than Bayes' Rule. Thus, the mean of a Bayesian-inferred distribution over states is the unique optimal estimate – for *any* strictly proper reward scheme. We formalize this in the following theorem:

Theorem 4. *If ρ is drawn from an ensemble $\pi_0(\rho)$, and a measurement $\mathcal{M}_j$ with conditional probability $p\left(\mathcal{M}_j|\rho\right)$ is observed, then every strictly proper scoring rule $\overline{R}(\rho : \sigma)$ is maximized by:*

1. *Using Bayes' Rule and Born's Rule:*

$$\pi_0(\rho) \longrightarrow \pi_{\mathcal{M}}(\rho) = \frac{p(\mathcal{M}|\rho)\pi_0(\rho)}{\int \mathrm{d}\rho \pi_0(\rho) p(\mathcal{M}|\rho)} = \frac{\left[\prod_{i=1}^N \mathrm{Tr}(E_i\rho)\right]\pi_0(\rho)}{\int \mathrm{d}\rho \left[\prod_{i=1}^N \mathrm{Tr}(E_i\rho)\pi_0(\rho)\right]}.$$

2. *Reporting the mean of* $\pi_{\mathcal{M}}(\rho)$.

This applies not only to relative entropy, our preferred measure of honesty, but to *any* honesty-guaranteeing reward scheme. We conclude that Bayesian inference is the unique solution to *honest* state estimation.

Other procedures will not optimize any measure of honesty derived from a strictly proper scoring rule. Alternative measures of honesty will either (a) in some circumstances, motivate an experimentalist to flat-out lie about the state, or (b) not be operationally implementable (e.g., fidelity). Our previous discussion of fidelity illustrates that a non-operational metric that guarantees the honesty of a *knowledgeable* experimentalist can fail dramatically in the face of uncertainty.

Information theorists have previously interpreted relative entropy as a good measure of two states' distinguishability [16, 17] – indeed, as the only meaningful one in the limit of many copies. We have invoked the Likelihood Principle rather than the many-copy limit, but we can easily allow Bob to jointly measure N copies of ρ. He must then apply a SPSR to the result, and Alice can expect a reward $\overline{R}(\rho^{\otimes N} : \sigma^{\otimes N})$. As $N \to \infty$, relative entropy remains meaningful, unlike other measures (e.g., the Brier score [10]).

In the presence of uncertainty, pure (or rank-deficient) states are infinitely dishonest estimates. Estimating a pure state means predicting that some event will *never* happen. Anyone taking such a prediction seriously would be justified in betting infinitely much money, at arbitrarily bad odds, against that event – and should therefore expect to lose infinitely much, if the estimate is incorrect. This should be discouraged.

Bayesian state estimation has been discussed previously [18, 19, 20], especially for pure states [21, 22]. The predominance of other methods such as maximum likelihoood, in the current literature (e.g., [23, 24, 25, 26, 27, 28, 29, 30], indicates that it has not received the attention it deserves. Our goal in this letter is to provide a concrete and compelling argument for Bayesian state estimation – and to call attention to the problematic implications of pure-state estimates.

Acknowledgments

RBK thanks Daniel James (for pointing out the problem in the first place), Jon Yard, John Preskill, and NSF contract PHY-0456720. PH thanks Howard Barnum, Simon Benjamin and especially Chris Fuchs for helpful discussions, as well as the Canada Research Chairs program, the Canadian Institute for Advanced Research and NSERC for support.

APPENDIX A: FURTHER RESTRICTING THE REWARD FUNCTION

We show that the quantum reward function is more tightly constrained than the classical one: the constants C_i of Theorem 1 must all be equal. We can assume that $D = 1$. Assuming a reward function of the form specified by Theorem 1, the inequality $\overline{R}(\rho : \sigma) \le \overline{R}(\rho : \rho)$ is then equivalent to

$$S\left(\rho\|\sigma\right) + \sum_i C_i(r_i - p_i) \ge 0. \tag{A1}$$

Let $U(t)$ be a smooth curve in the unitary group defined on a neighborhood of $t = 0$ and such that $U(0) = I$. Also let $\sigma(t) = U(t)\rho U^\dagger(t)$ and

$$g(t) = S\left(\rho\|\sigma(t)\right) + \sum_i C_i\big[r_i - \mathrm{Tr}\big(\rho U(t)|e_i\rangle\langle e_i|U^\dagger(t)\big)\big]$$

be the function defined by substituting $\sigma(t)$ into the expression of Eq. (A1). (Here $\rho = \sum_i r_i |e_i\rangle\langle e_i|$, so that $|f_i\rangle = U\,|e_i\rangle$.) Differentiating gives

$$\frac{d}{dt} S\left(\rho\|\sigma(t)\right) = -\mathrm{Tr}\big[\rho\big(\dot{U}(\ln\rho)U^\dagger + U(\ln\rho)\dot{U}^\dagger\big)\big], \tag{A2}$$

where $\dot{U} = dU/dt$ and the dependence of U on t has been suppressed. Because $\dot{U} + \dot{U}^\dagger = 0$ when $t = 0$, $\frac{d}{dt}S\left(\rho\|\sigma(t)\right)|_{t=0} = 0$. Likewise, $\dot{p}_i$ equals 0 when $t = 0$. We must therefore determine the second derivative of g at $t = 0$ and show that for a suitable choice of curve, this derivative is negative. In that case, $g(t) = g(0) + \ddot{g}(0)t^2/2 + O(t^3)$ with $\ddot{g}(0) < 0$, implying that $g(t)$ is negative for sufficiently small t.

So, differentiating again, we find

$$\frac{d^2}{dt^2} S\left(\rho\|\sigma(t)\right)\Big|_{t=0} = 2\,\mathrm{Tr}\Big[[X,\rho](\ln\rho)X\Big], \tag{A3}$$

where X is a Hermitian matrix such that $\dot{U} = iX$. We have made use of the identity $\ddot{U} + \ddot{U}^\dagger = -2\dot{U}\dot{U}^\dagger$ that can be proved by differentiating $UU^\dagger = I$. Because $S\left(\rho\|\sigma(t)\right) \ge 0$, the expression in Eq. (A3) must also be nonnegative.

Differentiating the second term of $g(t)$, we find that

$$\frac{d^2 g}{dt^2}\Big|_{t=0} = 2\,\mathrm{Tr}\Big[[X,\rho](\ln\rho + B)X\Big], \tag{A4}$$

where $B = \sum_i C_i |e_i\rangle\langle e_i|$.

We will now show that all the C_i must be equal. Assume without loss of generality that $C_1 \neq C_2$. Let $r_j = \exp(-(C_j + 2\ln 2)/2)$ for $j = 1, 2$ and $r_3 = 1 - r_1 - r_2$.

With these choices, $r_1, r_2 \leq 1/2$, making ρ a density operator. Now write $\tilde{\rho}$ and $\tilde{B}$ for the restriction of ρ and B to the span of the first two eigenvectors of ρ. Observe that there exists a choice of X also with support only on this subspace such that $\mathrm{Tr}\big[[X,\tilde{\rho}](\ln\tilde{\rho})X\big] > 0$.

With these choices, and noting that $\mathrm{Tr}\big[[X,\tilde{\rho}]X\big] = 0$,

$$\ddot{g}(0) \;=\; -2\,\mathrm{Tr}\Big[[X,\tilde{\rho}](\ln\tilde{\rho})X\Big] < 0. \qquad \text{(A5)}$$

Because $\ddot{g}(0) < 0$ contradicts the requirement that $D(\rho\|\sigma(t)) \geq 0$, we conclude that $C_i = C_1$ for all i. □

[1] K. Vogel and H. Risken, Phys. Rev. A **40**, 2847 (1989).

[2] G. M. D'Ariano, M. G. A. Paris, and M. F. Sacchi, Advances in Imaging and Electron Physics **128**, 205 (2003).

[3] Z. Hradil, Phys. Rev. A **55**, R1561 (1997).

[4] Z. Hradil, J. Rehcek, J. Fiursek, and M. Jezcaronek, in *Quantum state estimation*, edited by M. G. A. Paris and J. Rehacek (Berlin, Germany : Springer-Verlag, 2004, 2004), vol. 649 of *Lecture Notes In Physics*, pp. 59–112.

[5] E. Bagan, A. Monras, and R. Muoz-Tapia, Phys. Rev. A **71**, 062318 (2005).

[6] F. Tanaka and F. Komaki, Phys. Rev. A **71**, 052323 (2005).

[7] J. Rehacek, B.-G. Englert, and D. Kaszlikowski, arxiv.org/quant-ph **05**, 0405084 (2004).

[8] J. Bernardo, Annals of Statistics **7**, 686 (1979).

[9] M. Ohya and D. Petz, *Quantum entropy and its use*, Texts and monographs in physics (Springer-Verlag, Berlin, 1993).

[10] T. Gneiting and A. E. Raftery, Tech. Rep. 463R, Dept. of Statistics, Univ. of Washington (2005).

[11] L. J. Savage, Journal of the American Statistical Association **66**, 783 (1971).

[12] R. Bhatia, *Matrix Analysis* (Springer, 1996).

[13] J. O. Berger and R. L. Wolpert, *The Likelihood Principle: A Review, Generalizations, and Statistical Implications (with discussion).*, IMS Lecture Notes-Monograph Series (Institute of Mathematical Statistics, 1988).

[14] J. Aczél, Journal of Mathematical Analysis and Applications **75**, 447 (1980).

[15] R. Jozsa, Journal of Modern Optics **41**, 2315 (1994).

[16] V. Vedral, M. B. Plenio, K. Jacobs, and P. L. Knight, Phys. Rev. A **56**, 4452 (1997).

[17] C. A. Fuchs, Ph.D. thesis, University of New Mexico (1995).

[18] R. Derka, V. Buzek, G. Adam, and P. L. Knight, Journal of Fine Mechanics and Optics **11-12**, 341 (1996).

[19] R. Schack, T. A. Brun, and C. M. Caves, Physical Review A **64**, 014305 (2001).

[20] F. Neri, arxiv.org/quant-ph p. 0508012 (2005).

[21] K. R. W. Jones, Annals of Physics **207**, 140 (1991).

[22] K. R. W. Jones, Physical Review A **50**, 3682 (1994).

[23] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Physical Review A **64**, 052312 (2001).

[24] K. J. Resch, P. Walther, and A. Zeilinger, Phys. Rev. Lett **94**, 070402 (2005).

[25] Y. S. Weinstein, T. F. Havel, J. Emerson, N. Boulant, M. Saraceno, S. Lloyd, and D. G. Cory, Journal of Chemical Physics **121**, 6117 (2004).

[26] J. L. O'Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White, Physical Review Letters **93**, 080502 (2004).

[27] S. H. Myrskog, J. K. Fox, M. W. Mitchel, and A. M. Steinberg, Phys. Rev. A **72**, 013615 (2005).

[28] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Hffner, W. Hnsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt, Phys. Rev. Lett **92**, 220402 (2004).

[29] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, Physical Review Letters **90**, 193601 (2003).

[30] H. Haeffner, W. Haensel, C. F. Roos, J. Benhelm, D. C. al kar, M. Chwalla, T. Koerber, U. D. Rapol, M. Riebe, P. O. Schmidt, et al., Nature **438**, 643 (2005).

[31] Characterizing Bob's *non*-projective measurements is more complicated. For now, we note that: (1) in general, a fixed informationally complete POVM suffices; (2) our results remain largely unchanged, but Alice's expected reward may involve $\mathcal{C}[\rho]$ and $\mathcal{C}[\sigma]$ instead of ρ, σ, where $\mathcal{C}$ is a quantum channel.

[32] Jozsa showed [15] that $F(\sigma,\rho)$ is maximal if and only if $\sigma = \rho$